arXiv:1506.09158v1 [cs.DS] 30 Jun 2015

ON FAIR SIZE-BASED SCHEDULING

MATTEO DELL'AMICO, DAMIANO CARRA, AND PIETRO MICHIARDI

Abstract. By executing jobs serially rather than in parallel, size-based scheduling policies can shorten time needed to complete jobs; however, major obstacles to their applicability are fairness guarantees and the fact that job sizes are rarely known exactly a-priori. Here, we introduce the Pri family of size-based scheduling policies; Pri simulates *any* reference scheduler and executes jobs in the order of their simulated completion: we show that these schedulers give strong fairness guarantees, since *no job completes later in* Pri *than in the reference policy.* In addition, we introduce PSBS, a practical implementation of such a scheduler: it works online (i.e., without needing knowledge of jobs submitted in the future), it has an efficient $O(\log n)$ implementation and it allows setting priorities to jobs. Most importantly, unlike earlier size-based policies, the performance of PSBS degrades gracefully with errors, leading to performances that are close to optimal in a variety of realistic use cases.

Matteo Dell'Amico is with Symantec Research Labs; this work was done while he was at EURECOM. Email: matteo_dellamico@symantec.com. Phone: +33 4 93 00 82 61. Address: Symantec Research Labs at EURECOM, Campus SophiaTech, 450 Route des Chappes, 06410 Biot, France.

Damiano Carra is with University of Verona. Email: damiano.carra@univr.it. Phone: +39 045 802 7059. Address: Strada Le Grazie 15, Verona, Italy.

Pietro Michiardi is with EURECOM. Email: pietro.michiardi@eurecom.fr. Phone: +33 4 93 00 81 45. Address: Campus SophiaTech, 450 Route des Chappes, 06410 Biot, France.

1. Introduction

Schedulers are often based on *fair sharing*, where the resources are divided among jobs according to some fairness concept. The simplest case is processor sharing (PS), which partitions the resources equally among pending jobs at every instant. However, if users care about job completion time, rather than instantaneous job progression, sharing is not optimal: this is shown in FSP [3], a scheduler that optimizes job completion times while providing strong fairness guarantees. FSP dominates PS, i.e., *no job will complete later in FSP than in PS*, and it is based on a simple idea: schedule the job that would complete first in PS. Here, we discuss two issues arising from implementing a scheduler inspired by FSP in a practical context [7].

First, what if our concept of fairness is more elaborate than simple equal sharing? Many real-world schedulers have a flexibility which goes even beyond that of priority classes: e.g., the Hadoop capacity scheduler [9] applies a hierarchical concept of fair sharing to guarantee resources to units within an organization. We thus introduce a generalization of FSP's dominance result: given *any* scheduler, simulating it and executing jobs one at a time according to the order in which they would complete dominates the scheduler itself. Therefore, *any* scheduler can be used as a reference for fairness, and executing jobs serially is always beneficial.

Second, what happens when job size is only known approximately? Indeed, in practical settings, job sizes are rarely known a-priori. We thus introduce our work on scheduling based on inexact sizes and show that, if a scheduler has been designed without considering that the information about job size may be inaccurate, estimation errors may have dramatic impact on the performance for different workload characteristics. On the other hand, if the consequences of the estimation errors are properly addressed – as we do in our proposal, PSBS [1] – the scheduler performs close to optimally in a variety of workloads. PSBS is efficient (its complexity is $O(\log n)$ compared to $O(n)$ of FSP), and allows job priorities.

We conclude by highlighting open questions and future research directions related to scheduling with inexact job sizes.

2. Dominance Results With Known Job Sizes

We consider here the single-machine scheduling problem with release times and preemption; our goal, that materializes in the Pri scheduler, is to minimize the sum of completion times (according to Graham et al. [4], the $1|r_i; pmtn|\sum C_i$ problem) with the additional dominance requirement: no job should complete later than in a scheduler which is taken as a reference for fairness. Without this limitation, the optimal solution is the Shortest Remaining Processing Time (SRPT) policy. We call *schedule* a function $\omega(i,t)$ that outputs the fraction of system resources allocated to job i at time t. For example, for the processor-sharing (PS) scheduler, when n jobs are *pending* (released and not yet completed), $\omega(i,t) = \frac{1}{n}$ if job i is pending and 0 otherwise. Furthermore, we call $C_{i,\omega}$ the completion time of job i under schedule ω.

Definition 1. Schedule ω *dominates* schedule ω' if $C_{i,\omega} \leq C_{i,\omega'}$ for each job i.

Our scheduler prioritizes jobs according to the order in which they complete in ω.

Definition 2. A *completion sequence* $S = [s_1, \ldots, s_n]$ is an ordering of the jobs to be scheduled. A schedule ω *has completion sequence* S if $C_{s_i,\omega} \leq C_{s_j,\omega} \forall i < j$.

Definition 3. For a completion sequence S, the Pri_S schedule is such that $\text{Pri}_S(i,t) = 1$ if i is the first pending job to appear in S; $\text{Pri}_S(i,t) = 0$ otherwise.

We now show that scheduling jobs in the order in which they complete under ω dominates ω.

Theorem. Pri_S *dominates any schedule with completion sequence* S.

Proof. We have to show that $C_{i,\text{Pri}_S} \leq C_{i,\omega}$ for each job i and any schedule ω with completion sequence S. Let j be the position of j in S (i.e., $i = s_j$); we call M the minimal makespan of the $S_{\leq j} = \{s_1, \ldots, s_j\}$ set of jobs,[1] and we show that $C_{i,\text{Pri}_S} \leq M$ and $M \leq C_{i,\omega}$:

[1]The *makespan* of a set of jobs is the maximum among their completion times, therefore $M = \min_{\omega \in \Omega} \max_{i \in \{1,\ldots,j\}} C_{S_i,\omega}$ where Ω is the set of all possible schedules.

- $C_{i,\mathrm{Pri}_S} \leq M$: minimizing the makespan of $S_{\leq j}$ is equivalent to solving the $1|r_i; pmtn|C_{\max}$ problem applied to the jobs in $S_{\leq j}$: this is guaranteed if all resources are assigned to jobs in $S_{\leq j}$ as long as any of them are pending [5]. Pri_S guarantees this, hence the makespan of $S_{\leq j}$ using Pri_S is M. Since $i \in S_{\leq j}$, $C_{i,\mathrm{Pri}_S} \leq M$.
- $M \leq C_{i,\omega}$ follows trivially from the fact that ω has completion sequence S and, therefore, $C_{i,\omega}$ is the makespan for $S_{\leq j}$ using schedule ω. □

This theorem generalizes the results by Friedman and Henderson [3]: FSP follows from applying Pri_S to the completion sequence of PS. The generalization is important: in practice, one can define a scheduler that provides a desired type of fairness, and then optimize the performance in terms of completion time by applying the Pri_S scheduler. For instance, assume that the system deals with different classes of jobs that have different weights, and the scheduler to apply to provide fairness is the discriminatory processor sharing (DPS): the theorem guarantees that Pri_S dominates DPS. We have exploited exactly this results in our scheduler PSBS [1], which, in the absence of errors, dominates DPS. Note that both FSP and PSBS are appliable *online*: even without information on future jobs, it is possible to compute which pending job completes first in PS and DPS and hence decide which job to schedule.

3. Scheduling With Approximated Sizes

The results above seem to suggest that size-based schedulers should be employed ubiquitously; however, a major obstacle to their applicability is that, in a large majority of cases, job size cannot be known exactly; on the other hand, it is often possible to compute an *estimate*. In this Section, we synthetize our results on the topic [1, 2].

Only a few other works [6, 8] tackle the problem of scheduling with inaccurate sizes; they show rather pessimistic results, suggesting that size-based schedulers outperform non size-based counterparts only when estimations are precise. We have complemented those works with an extensive simulative study, generating synthetic workloads with several varying parameters related to the job size distribution and the error distribution. To help reproducibility, our simulator is available as free software.[2] We have found that job size distribution plays an essential role: if job sizes are skewed (i.e., a few very large jobs make up a large fraction of the total work), size estimation errors cause serious performance issues in existing size-based schedulers. This phenomenon is mainly caused by the fact that, if the size of a large job is underestimated, it gains positions in the queue; when it enters in the service, it reaches a point when it cannot be preempted and it blocks the server until it has completed (at the expense of jobs that are actually small). The opposite situation, a job with overestimated size, instead, has little impact on the other jobs (see [1] for an illustrative example).

Our proposal, PSBS, leverages these last observations. The main idea is to react when the system detects that a job has been underestimated, i.e., when a job is taking more resources than the ones initially estimated – we call these jobs *late*. The scheduler treats the late jobs differently, and lets other jobs be served, so that the impact of the underestimation is limited.

This, coupled with the fact that PSBS schedules jobs one at a time according to their completion time when simulating DPS, allows PSBS to obtain close to optimal performance, yet guaranteeing the same fairness of the simulated scheduling (DPS). PSBS has been inspired by FSP but, unlike FSP, it allows setting priorities and it deals with estimation errors. Moreover, PSBS is efficient, since its complexity is $O(\log n)$, compared to $O(n)$ of FSP. Note that, by properly configuring its parameters and with no estimation errors, PSBS behaves as FSP, therefore our efficient implementation represents a gain with respect to FSP.

In Figure 3.1, we show the ratio between the mean sojourn time (MST)[3] obtained using size-based schedulers, such as SRPT and FSP, normalized against the MST of PS. Job priorities are homogeneous and sizes are generated according to a Weibull distribution (heavy-tailed for shape < 1, light-tailed otherwise); the relative job size estimation error is distributed according to a

[2] `http://github.com/bigfootproject/schedsim`

[3] A job's *sojourn time* is the interval between its release and completion; minimizing MST also minimizes $\sum C_i$.

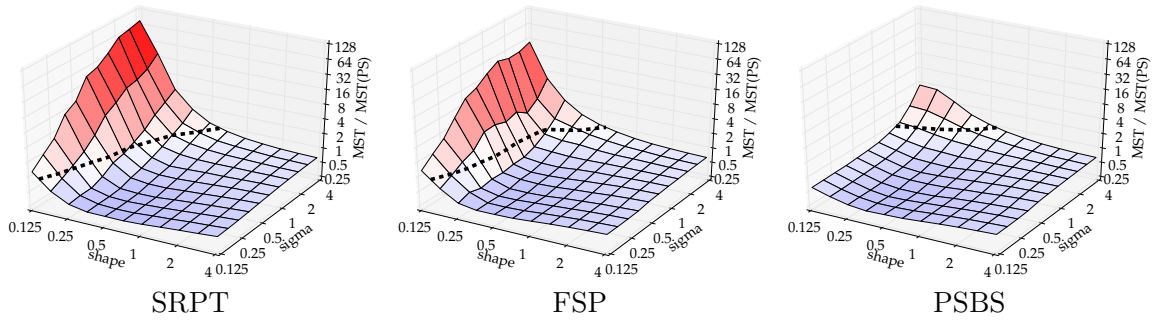

FIGURE 3.1. Mean sojourn time using size-based schedulers, normalized against PS.

log-normal distribution – i.e., higher values of sigma yield larger errors. Job inter-arrival times are distributed exponentially and the load (ratio between arrival and service rate) is set to 0.9.

The results of Figure 3.1 show that SRPT and FSP suffer when the workload is highly skewed, even with moderate estimation errors; on the other hand, PSBS largely corrects this issue, and it is outperformed by PS only in extreme cases where the workload is very skewed *and* job size estimation is very imprecise (sigma greater than 2, which corresponds to a correlation coefficient between the job size and its estimate less than 0.15). In fact, PSBS performs close to optimally in most cases; similar results are obtained when playing back real workloads on the simulator. Other parameters such as load and job inter-arrival times do not have a large impact on the results.

4. Conclusion

We have shown two main results. For the first time in this paper, we have generalized the dominance results of FSP over PS to *any* "simulated" scheduler. We also have synthetized our results on scheduling based on approximate job sizes: PSBS largely mitigates the problem of existing policies, and often obtains close to optimal results, yet maintaining the desired fairness among jobs. We consider that these results can be interesting both for theorists and for practitioners. In practical systems, the PSBS policy can be used as a basis to build efficient size-based schedulers, as our work for Hadoop [7] demonstrates.

With respect to theory, we identify a problem that may be of interest to the community. We have shown that a scheduler such as PSBS can perform well in a variety of realistic use cases; we are able to explain these results with intuition and substantiate them with numerical experiments. However, an analytical characterization of the problem may provide more insights on the situations where such a strategy performs well, and a way to predict scheduler performance in function of the workload characteristics. Such a modeling approach could be useful to go beyond PSBS. We speculate that job size information – even if approximated – is better than no information, and hence we conjecture that it is possible to design a scheduler that *always* outperforms non size-based counterparts, when the distribution of estimation errors is known.

References

[1] M. Dell'Amico et al. PSBS: Practical size-based scheduling. *arXiv:1410.6122*, 2014.
[2] M. Dell'Amico et al. Revisiting size-based scheduling with estimated job sizes. In *MASCOTS*. IEEE, 2014.
[3] E. J. Friedman and S. G. Henderson. Fairness and efficiency in web server protocols. In *ACM SIGMETRICS Performance Evaluation Review*, volume 31, pages 229–237. ACM, 2003.
[4] R. L. Graham et al. Optimization and approximation in deterministic sequencing and scheduling: a survey. *Annals of discrete Mathematics*, 5:287–326, 1979.
[5] C. L. Liu and J. W. Layland. Scheduling algorithms for multiprogramming in a hard-real-time environment. *JACM*, 20(1):46–61, 1973.
[6] D. Lu et al. Size-based scheduling policies with inaccurate scheduling information. In *MASCOTS*. IEEE, 2004.
[7] M. Pastorelli et al. HFSP: size-based scheduling for Hadoop. In *BIGDATA*. IEEE, 2013.
[8] A. Wierman and M. Nuyens. Scheduling despite inexact job-size information. In *ACM SIGMETRICS Performance Evaluation Review*, volume 36, pages 25–36. ACM, 2008.
[9] M. Zaharia. Job scheduling with the fair and capacity schedulers. In *Hadoop Summit*, 2009.